

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 May 2002 (30.05.2002)

PCT

(10) International Publication Number
WO 02/43004 A2

(51) International Patent Classification⁷: G06T 7/00

(21) International Application Number: PCT/GB01/05094

(22) International Filing Date:
19 November 2001 (19.11.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0028491.9 22 November 2000 (22.11.2000) GB

(71) Applicant (for all designated States except US): ISIS IN-
NOVATION LIMITED [GB/GB]; Ewert House, Ewert
Place, Summertown, Oxford OX2 7SG (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): MULET-PARADA,
Miguel [ES/GB]; St Cross College, Oxford, Oxfordshire
OX1 3LZ (GB). NOBLE, Julia, Alison [GB/GB]; Dept. of
Engineering Science, Oxford, Oxfordshire OX1 3PJ (GB).

(74) Agents: NICHOLLS, Michael, John et al.; J.A. Kemp &
Co., 14 South Square, Gray's Inn, London WC1R 5JJ (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN,
YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.



WO 02/43004 A2

(54) Title: DETECTION OF FEATURES IN IMAGES

(57) Abstract: An image processing technique which identifies pixels in images which are associated with features having a selected shape, such as but not exclusively step edge, roof, ridge or valley. The shape of the intensity profile in the image is compared in an intensity independent way with a shape model to select those pixels which satisfy the shape model and are thus associated with the feature of interest. This comparison is achieved by examining the phase and amplitude of a spectral decomposition of parts of the image profile in the spatial or spatio temporal frequency domain. This decomposition can be achieved using quadrature wavelet pairs such as log-Gabor wavelets. The difference between the odd and even components, known as the feature asymmetry, gives an indication of the shape of the feature. The analysis may be extended to the time domain by looking at the shape of the image profile across a time sequence of images, which gives an indication of the velocity of a moving feature. Pixels identified as belonging to a feature of the right shape are labelled with the value of feature asymmetry, the local amplitude, feature orientation and feature velocity, and this information can be used to improve the tracking of detected features through a sequence of images.

DETECTION OF FEATURES IN IMAGES

The present invention relates to a method and apparatus for processing images to detect particular features in the image. It is particularly concerned with the detection
5 of features in noisy 2D, 2D + T, 3D and 3D + T images, and is thus particularly useful in medical imaging, for instance using ultrasound.

There are a variety of techniques for processing images to enhance them, particularly to reduce the amount of noise and/or to detect features of interest in the image. Such
10 techniques are particularly useful in areas where images are noisy, for instance in medical imaging. For example, ultrasound images include a considerable amount of speckle which makes images difficult to interpret. Skilled operators are able to interpret such images based on their experience and knowledge of the features of interest likely to be present. However, automating the process of detecting features of
15 interest is made difficult by the presence of noise and imaging artefacts. Similar problems arise in other imaging techniques, such as magnetic resonance imaging.

In general up to now techniques proposed for enhancing images have been based on examining the intensity of the image, for instance by smoothing the intensity to
20 reduce noise or by setting a detection threshold based on amplitude. However, techniques based on thresholding are not entirely satisfactory because of the difficulty of setting a suitable global threshold. For instance, ultrasound images include attenuation artifacts caused by the changing orientation and reflectivity of the tissue being imaged. Choosing one threshold which is suitable for excluding noise, but
25 including all features of interest, is difficult if not impossible.

Other techniques have been proposed for analysing noisy images for detecting features which are based on correlating images over time or with movement of the imaging probe. For instance, ultrasound speckle decorrelates with movement of the
30 probe and over time. A problem with some of these techniques, however, is that they tend to blur the features of interest in the image. Thus the user has to accept to trade-

-2-

off between removing noise and losing image quality.

It is therefore an object of the invention to provide a technique for automatic detection of features of interest in an image which improves on the prior art techniques.

5

The present invention provides a method for detecting features of interest in an image based on the shape of the intensity profile of the feature, rather than its intensity. This is achieved by an intensity independent comparison of the intensity profile across the image with a shape model of features of interest. To improve subsequent processing of the image, areas of the image which are detected as corresponding to the shape model are labelled with a description of their properties.

10

In more detail, therefore, the present invention provides a method of detecting features-of-interest in an image, comprising the steps of:

15

making an intensity independent comparison of the shape of regions of the intensity profile across the image with a shape model for predefined features of interest to define as features of interest areas of the image satisfying the shape model; and

20

storing for each of said defined areas a label comprising the level of similarity to the shape model.

The predefined features of interest may be positive and negative step edges, or roof or ridge shapes, or valleys or other features as desired depending on the application.

25

The label may include a measure of the orientation of the feature in the image, and the comparison may also be performed on an intensity profile taken across successive frames of an image sequence to derive a measure of the velocity of image features. The label may then comprise a measure of the velocity.

30

The comparison with the shape model is advantageously performed in the spatial or spatio temporal frequency domains, thus by decomposing the intensity profile into

-3-

spatial frequency components and examining the phase and amplitude of those components. In one embodiment this is achieved by convolving the intensity profile with a quadrature filter to derive the phase and amplitude of a pair of odd and even components of the profile in the spatial frequency domain. The phase and amplitude
5 of these components is characteristic of different shapes in the intensity profile. The difference between the odd and even components, which is a measure of the "feature asymmetry" is, for example, a maximum for a step edge. Thus by setting constraints on the value of the feature asymmetry it is possible to determine that the area of the image under consideration corresponds to the shape being sought. The values of
10 feature asymmetry and local amplitude (which is based on the amplitude of the two components), are included in the label.

Preferably the filters are quadrature wavelet filters, e.g. log-Gabor filters, so that the intensity profile is convolved with odd and even wavelets. The scale of the wavelet is
15 chosen in accordance with the scale of the features to be detected, and to be much larger than the scale of noise in the image. This means that the technique selects features according to their shape and scale, but regardless of the value of intensity.

The filters may be oriented in different directions across the image to detect image
20 features having different orientations. The label may then include a measure of the feature orientation, which can be derived from the relative responses of the differently oriented filters.

The image processing technique of the invention may be used as a precursor to
25 tracking detected image features through a sequence of image frames. The provision of the labels carrying information about the detected features is particularly useful in such a tracking process. For instance, rather than simply searching for the closest image feature in two successive frames, it is possible also to compare the labels of the detected image features in the two frames and to define them as relating to the same
30 image feature if the labels satisfy predefined search criteria.

Such search criteria may include a condition on the value of the feature asymmetry of the two features being detected and a condition on the orientation of the detected features (for instance that they are oriented similarly).

- 5 Where the label includes a measure of the velocity of the feature, the search area from frame-to-frame may be defined in accordance with the velocity, and the search criteria may include a condition on the velocity.

10 The invention is particularly applicable to ultrasound or magnetic resonance images, but also to x-ray and other imaging modalities. The embodiment described below relates to echocardiography, but the invention is applicable to ultrasound images in general, such as of different organs and tissues e.g. coronary arteries, liver, foetus etc, and to different ultrasound modalities such as the use of contrast and different signal processing techniques such as Doppler imaging and harmonic imaging. For instance
15 in echocardiography, by setting the shape model appropriately the invention is adaptable to the detection of any features corresponding to a step change in intensity, such as the ventricular walls.

The invention may be embodied in a system adapted to perform the image processing
20 method, and it may be embodied by a computer program comprising program code means for executing the method. Thus the invention further provides for a computer program storage medium carrying such a computer program, and also a computer system programmed to carry out the method.

- 25 The invention will be further described by way of example with reference to the accompanying drawings in which:-

Figures 1(a) and (b) illustrate Fourier decompositions of a step edge and a triangular function;

30

Figure 2 illustrates the values of the local phase signatures of spatial

-5-

frequency components for different functions;

Figures 3 (a) and (b) illustrate respectively even and odd two-octave log-Gabor wavelet functions;

5

Figure 4 illustrates the spectral properties of the log-Gabor wavelets;

Figure 5 (a) illustrates a typical intensity profile across an ultrasound image;

10

Figure 5 (b) illustrates the local amplitude result of a convolution of an intensity profile with log-Gabor filters at a succession of scales;

Figure 5 (c) illustrates the local phase result of a convolution of an intensity profile with log-Gabor filters at a succession of scales;

15

Figure 6 schematically illustrates an image processing system according to an embodiment of the invention;

20

Figures 7 (A) and (B) illustrate examples of the feature orientations in two and three dimensions used in the embodiment of Figure 6;

Figure 8 illustrates the effect on feature detection of analysing the intensity profile using log-Gabor wavelets of different scale (wavelengths);

25

Figure 9 illustrates a process for tracking a detected feature through a sequence of images;

Figure 10 schematically illustrates part of the tracking process; and

30

Figure 11 illustrates the relationship between local phase and feature asymmetry.

-6-

While traditional methods for detection of a step edge in an intensity profile rely on examining the derivatives of the intensity profile (i.e. its slope), to detect the sudden changes in intensity which mark the visible edge of features in the image, this embodiment of the invention examines a decomposition of the intensity profile in the spatial or spatio temporal frequency domains. For instance, Figure 1 illustrates Fourier decompositions of a step edge (Figure 1(a)) and a triangular function (Figure 1(b)). In Figure 1 (a) it can be seen that three sinusoidal components i, ii and iii sum to make an approximation to a step edge iv. Further it will be noted that the phase of all three components is the same, zero, at the positive-going step edge. Further the phase value of all the components at the negative step edge is 180 degrees. Reference to Figure 1 (b) illustrates that for a triangular peak the phase value of all of the Fourier components is 90 degrees at the peak of the triangle.

The "local phase" is defined as the phase offset of each Fourier component describing a signal, measured at the event of interest. For features such as a step, the peak of a roof or the foot of a ramp, the local phase of all of the Fourier components converges to the same value. Further, the value is characteristic of the type of feature. The phase values for different types of feature is illustrated in the diagram of Figure 2. It can be seen, therefore, that the phase of the spectral components of a signal gives an indication of the shape of the signal. It should be noted that the particular value of phase is arbitrary (depending on the position of the origin), what is important is that the different features have different, known values.

While Fourier decomposition might be suitable for perfectly shaped features, the intensity profile across typical images is much more complex and noisy and a Fourier decomposition of such a complex waveform would be extremely complex. Instead, this embodiment of the invention looks in turn at successive small sections of the intensity profile across the image by making the spectral analysis using wavelets. This embodiment uses log-Gabor wavelets as illustrated in Figure 3 (which are in fact Gaussian modulated sinusoids) which have a non-zero central section and are zero elsewhere. Thus convolving them with the intensity profile effectively windows a

-7-

small region of the profile. Figure 3 (a) shows the even wavelet based on a Gaussian modulated cosine, and Figure 3 (b) the odd wavelet based on a Gaussian modulated sine function. The spectral properties (plotted against frequency and log-frequency) of the wavelets are shown in Figure 4.

5

Thus just as a Fourier analysis of a profile gives the amount and phase of each of the Fourier components, convolving the intensity profile with the log-Gabor wavelet pair extracts the amplitude and phase of the odd and even components for the region of interest.

10

The scale of the wavelets used is important. Figure 5 (a) illustrates a typical intensity profile across an echocardiographic image. One can imagine analysing this profile using very small scale versions of the wavelets shown in Figures 3 (a) and 3 (b). This would result in an analysis of the very small scale variations in intensity in the profile.

15

On the other hand, if the scale of the wavelets used was much larger, such that, for example, only one wavelet fitted across the entire length of the profile, the effect of the small scale variations in intensity would not affect the result, and the analysis would only actually pick up the broad overall shape of the profile. Thus the effect of varying the scale of the wavelet is to look for different scale features in the profile.

20

The effect of varying the wavelength, or scale, of the wavelet is illustrated by the scalograms of Figures 5 (b) and (c). These scalograms are constructed by taking the convolutions of a sequence of log-Gabor filters of increasing wavelength, and stacking vertically the output local phase or local amplitude responses of each filter, so that each forms a new row of the image. The scalograms on the left (Fig. 5B) are computed from the local amplitude responses given by

25

$$A(x) = \sqrt{o^2(x) + e^2(x)}$$

30

where $e(x)$ is the output of the even log-Gabor wavelet, and $o(x)$ is its odd counterpart. Similarly, the phase scalogram of Fig. 5B is computed from the absolute

-8-

local phase values given by:

$$\Phi(x) = \left\| \arctan \left(\frac{e(x)}{o(x)} \right) \right\|$$

On the left scalogram, the intensity values correspond to the output amplitudes of the convolutions. The phase values on the right have been mapped to intensities so that black corresponds to zero phase (positive steps) and white corresponds to π radians (negative steps). Peaks, having $\pi/2$ phase signatures, are grey.

The phase scalogram shows that at small scales many features corresponding to these extreme values are found. However, as the filter scale increases, the number of features reduces to a few straight bands. Among these it is possible to identify the step edges (local phase converges to zero, thus black on the scalogram) corresponding to the epicardial (1) and the endocardial (2,3) step edges, as well as a pericardial peak (4). From the phase scalogram such as Fig. 5c, two observations can be made: 1) features such as step edges and peaks have *stable* phase values and 2) phase signatures remain *localised* over scales. Stability means that there is always a particular *phase signature* associated with each feature shape over scales (in the illustration the zero degree and π radian values corresponding to positive and negative steps respectively are shown). Localisation means that the local phase signature associated with a feature remains in the proximity of the event (i.e. at the same place on the x-axis) over a range of scales. Thus because localisation and stability hold, it is possible to identify features of interest from the phase output of a large scale convolution according to pre-defined phase signatures, and trace these back along straight paths to their location. Thus the system can use a large scale filter to remove noise and still identify feature points accurately, and by examining a phase scalogram such as Figure 5 (c) it is possible to set a suitable scale at which the phase signatures of features of interest are clear.

In the present embodiment of the invention a single scale is sufficient to detect echo

boundaries of interest corresponding to the heart. In different applications, such as viewing different organs or using different modalities it may be necessary to combine the output of the filters at different scales to enhance the detection. If necessary filters may be run at different scales and the results of different scales combined by, for
5 example, averaging feature asymmetry results and recording selected features if their shape description is consistent over a number of scales.

In this embodiment of the present invention the images being processed are echocardiograms of the left ventricle obtained using an HP-Sonos 5500 scanner
10 producing images of a size 300 x 250 pixels at a frame rate of 50 to 60 Hz. For these images scales of between 32 to 56 pixel are used with log-Gabor wavelets of two octaves.

So far, the issue of boundary detection has been centred on the 1D problem and the
15 analysis of a 1D profile. However, the present invention is also designed to detect and find the orientation of features (steps, valleys, ridges) in 2D image and 3D volume data (including 2D+T i.e. sequence data).

To do this a filter bank of oriented filters is designed to sample the space of the data:
20 i.e. a plane for a 2D image or a volume for 3D or 2D+T data. This embodiment of the invention has been applied to echocardiographic sequences, and in this case a 2D+T volume is created by stacking the successive image frames behind each other creating a 3D volume (such that the two spatial dimensions follow the x and y axes and the time dimension follows the z axis). A volume for three spatial dimensions (3D) is
25 created by stacking images taken at different depths.

Sampling of multidimensional data and detection of feature orientation require:

1. the extension of log-Gabor filters to higher dimensions (2D and 3D),
2. the choice of optimal orientations for the filters in the 2D and 3D spaces, and
30 3. an interpolation method that allows estimation of the orientation of the features detected from a small number of filter responses.

-10-

Extending a 1D filter function to higher dimensions is called 'spreading'. In this embodiment the log-Gabor filters are spread to higher dimensions by multiplying with a Gaussian function in the spectral domain. For example in 2D

$$G(\omega_r, \phi) = \exp - \left(\frac{(\log(\omega_r / \omega_{r0}))^2}{2(\log(\kappa / \omega_{r0}))^2} + \frac{(\phi - \phi_0)^2}{2\sigma_\phi^2} \right)$$

- 5 Here σ_ϕ defines the extent of the spreading function as a scaling s of the separation between filters $\sigma_\phi = s \times \Delta \phi$. This spreading concept can, of course, be extended to 3D.

- 10 The number of filters and their particular orientations have to be defined such that all possible features of interest are detected and their direction estimated reliably. In the absence of prior knowledge on the expected orientation of the features, this requires evenly sampling the data space with sufficient filters, so that an interpolation procedure can estimate the orientation of the feature accurately. In this embodiment three evenly spaced filters in 2D and six evenly spaced filters in 2D+T provide good
15 accuracy while minimising the number of filtering operations required. These are illustrated in Figures 7(A) and 7(B) respectively.

- Various interpolation algorithms can be used; in this embodiment, the filter responses are interpolated by using an elliptic (2D) or ellipsoidal (3D) fit to the filter responses.
20 This approximation provides estimates of feature orientation within the accuracy required by the tracking algorithm.

- Estimation of feature orientation in a 2D+T volume is equivalent to finding the orientation and direction of movement of a feature. If one imagines a horizontal line
25 moving upwards over a sequence of images, then it will appear in the frame stack as a plane with a given inclination to the horizontal proportional to the speed of the line. A fast line will become a steep plane; a slower line would be less steep. Estimating this component of feature orientation is equivalent to estimating feature velocity.

-11-

This can be used to anticipate the motion of anatomical boundaries and is of great benefit for subsequent tracking stages.

This embodiment of the invention is particularly applicable to the detection of anatomical boundaries such as heart chamber boundaries and in particular ventricular boundaries in the echocardiographic image of the left ventricle. These are step edges in intensity. In this embodiment such step edges are detected by deriving the "feature asymmetry" which is a measure of the difference between the odd and even filter outputs at each point in the image and is a function of the phase signature. Thus rather than looking at the value of the local phase as defined above, it is the feature asymmetry which is derived and this can be defined as:-

$$FA(\mathbf{x}) = \frac{|o(\mathbf{x})| - |e(\mathbf{x})|}{\sqrt{e^2(\mathbf{x}) + o^2(\mathbf{x})}}$$

where $o(x)$ and $e(x)$ are the odd and even filter outputs at point x in the 2D image or 2D+T volume.

The feature asymmetry is directly related to the local phase in the following way:-

$$FA(\mathbf{x}) = |\sin(\Phi(\mathbf{x}))| - |\cos(\Phi(\mathbf{x}))|$$

where Φ is the local phase.

This function is plotted in Figure 11. It can be seen that the feature asymmetry varies almost linearly with local phase ranging from -1 for peaks and valleys to +1 for positive and negative steps. (It should be noted that the phase values for Fig. 11 and the embodiment below are offset by 90° from those shown in Figs. 1 and 2. If the zero phase point had been defined as in Figures 1 and 2, then Fig. 11 would simply be

-12-

inverted, with $FA = -1$ for positive and negative steps (at 0° and 180°) and $+1$ for peaks and valleys (at 90° and 270°).)

5 With this embodiment of the present invention, therefore, it is the feature asymmetry which is calculated and thresholded to detect positive and negative steps. For instance, a threshold of 0.9 for the feature asymmetry corresponds to a local phase range of ± 5 degrees around ± 90 degrees and will thus detect positive and negative steps. The exact value of the threshold affects the width of the detected features. The threshold is dimensionless and thus can be applied to any image to detect the edges
10 that fall within the scale of the filters. Noise rejection is thus solely a matter of scale and feature shape, not intensity. Thus by deciding whether or not image points belong to a particular feature on the basis of feature asymmetry means that the technique is independent of the intensity of the image points. Of course the intensity can be noted and used in later processes, such as tracking as explained below.

15 Figure 8 illustrates the results of processing an echocardiogram of the left ventricle with the log-Gabor functions of Figure 3 at a variety of scales from 16 pixels through to 72 pixels. The white overlay corresponds to detected step edges and marks those pixels with feature asymmetry values greater than the threshold 0.9. It can be seen
20 that at the finest scale, for instance 16 pixels, most of the detected values correspond to the speckle texture in the ultrasound image. However, as the scale is increased it is the features of interest in the image which are detected, such as ventricular walls, for instance at 48 pixels scale. Thus for this image using a scale between 32 and 56 pixels would give satisfactory results. The actual scale to be used in any situation
25 depends on the image (for instance the quality of the image acquisition and the imaging depth setting used by the sonographer), but the user can be given control of the scale giving the ability to tune or focus the detection to obtain good delineation of the desired features of interest.

30 Figure 6 is an example that illustrates schematically the image analysis processing of this embodiment. The data 1 (which can be an image or sequence of images) is input

-13-

and Fourier transformed at 3 before being convolved with the series of oriented quadrature filters 5. The output of the quadrature filters 5 is then used to compute the feature shape measure, which in this embodiment is the feature asymmetry (which is a maximum for step edges). The feature asymmetry is derived from the even and odd filter outputs 7 using the formula above. The feature asymmetry is compared to a threshold (for instance 0.9) which is the only parameter used to accept image pixels as feature candidates via AND gates 11 for subsequent processing. For processing in two spatial dimensions three filters are required, for processing in 3D six filters are needed as shown in Figures 7(A) and (B). In Figure 7(A) the orientations are directed to the vertices of a regular hexagon. In Figure 7(B) the orientations are in the directions of the vertices of a regular icosahedron. For each of these candidate pixels an intensity-based measure, in this case the local amplitude as defined above, is computed at 9. Thus for each pixel accepted according to the feature asymmetry threshold, the local amplitude and feature asymmetry value from each of the oriented filters is recorded.

It will be appreciated that the value of the local amplitude from any of the filters depends on the relative orientation of the filter with respect to the feature. The maximum response is obtained when the filter is incident normal to the direction of the feature. Conversely, the minimum response is obtained when the filter is oriented parallel to the feature. In between, the local amplitude response decays smoothly and monotonically. This can be understood intuitively because the intensity profile along a line parallel to a step edge is constant, whereas the intensity profile along a line normal to the step edge gives the full step edge change. Thus the interpolator 13 determines the orientation of the feature from the local amplitude responses of the oriented filters 5. There are a variety of ways in which this can be done, such as a linear combination of the signals from the different orientations or by considering the outputs of the filters as vectors having a direction corresponding to the filter orientation and a length corresponding to the local amplitude, and fitting an ellipse to the end points of the vectors, the major axis of the fitted ellipse corresponding to the normal to the feature direction. The direction of the normal does not, of course,

-14-

distinguish between positive and negative edges, but these are distinguishable by the sign of the odd filter output.

5 A non-maximum suppression algorithm 14 can be run on the results to narrow the detected edges to approximately pixel wide contour segments. Further, an intensity-related noise rejection, and preferably thresholding and chain following, stage 16 is used to avoid streaking and to enhance the continuity of the contour features. The result is the output at 18 of a series of connected points which are detected as belonging to an edge feature, with a probability proportional to their local amplitude
10 and feature asymmetry scores, and labelled with an estimate of the feature orientation. In the case of sequences of images, an additional estimate of the feature plane with respect to the temporal axis gives an estimate of feature image velocity through the sequence.

15 Thus having detected pixels which belong to a particular type of image feature (i.e. a feature defined by a particular shape in the intensity profile), and having labelled those pixels with the local amplitude, feature asymmetry, orientation and optionally velocity, the system has provided a particularly rich description of the properties of the detected features. This rich description can be used to advantage in further
20 processing of the image, particularly when compared with techniques which just mark a pixel as belonging to a detected feature or not. An example of such further processing will be described below, with reference to Figures 9 and 10.

25 In the particular application considered here, to echocardiographic images of the left ventricle, the feature being detected is the left ventricular wall (endocardium) which is a closed contour having a shape as illustrated at 20 in Figure 9. As the heart beats the ventricle contracts as the ventricular wall moves and the contour detected in a sequence of images of the beating heart should follow the movement of the endocardium. While ideally it would be possible simply to detect the image features
30 corresponding to the endocardium in each frame of the sequence, in practice noise and artifacts tend to produce a disconnected representation of the cardiac boundaries,

-15-

and a number of spurious responses unrelated to the cardiac walls. Tracking the feature through a sequence of images can interpolate the features detected in the individual frames to produce a continuous contour which follows, or tracks, the feature in the image.

5

Various tracking techniques have been proposed, but tracking can be improved by using the rich description of the features provided by the detection technique above. This is illustrated in Figure 9. Firstly, taking a contour 20 inherited from an adjacent time frame, a search slab 22 is projected inwardly and outwardly from each point on the contour. Figure 9 illustrates an example in which these slabs are three pixels wide and twelve pixels long. Then each tangential row of three pixels within the slab is processed separately to find the strongest, local amplitude response of the three pixels, that agrees most closely with the orientation of the contour point and has the highest feature asymmetry score. This is repeated for each row of three pixels in the slab. In this embodiment one of the three pixels was selected as a possible candidate if it has the strongest local amplitude response amongst the three and has a feature asymmetry score greater than 0.8 and if the difference in orientation between the contour normal and the feature normal is less than 5 degrees. Thus the candidates are selected making use of the labels attached during the feature detection process. It should be noted that the orientation measure takes into consideration the direction of the step edge (whether it is a positive or negative step).

10
15
20

It should be noted that in the case of processing an image in two spatial dimensions plus time, the velocity estimate for the points can be used to restrict the search space along the normal. Thus the length of the search slab is selected according to the velocity of the feature.

25

The result of this is a one-pixel-wide slab 24 normal to the contour which consists of that pixel (if any) from each of the tangential rows in the search slab which satisfies the search criteria.

30

-16-

The next step is to draw the new contour using the results of the one-pixel wide normals 24. In this embodiment this is achieved by a dynamic programming technique which finds the optimal path linking a single point from each of the normals 24, the optimal path being defined with reference to a cost function. The cost function defines each transition between candidate points from one normal 24 to the next as having a cost associated with it. Thus, referring to Figure 10, the normals 24 are stacked next to each other and for a given pixel k in the ith row, the cost associated with a transition to the corresponding pixel in the next row or each of its neighbours is calculated.

10

The cost function is designed to blend the local amplitude values and orientation information and it includes partial costs such as:- a large penalty cost if an entry on the normal 24 has been left blank because no feature has been found in the image; the cost related to the local amplitude measurement, the distance away from the previous contour point and a smoothing term. The optimal path is the one that links one point from each normal such that the sum of all costs is minimal. A least squares transformation is then used to fit the contour inherited from the adjacent frame to the points selected by the dynamic programming routine.

15

Of course other techniques for tracking the contour can be used, but the aim is to improve the tracking process using the labels attached during the feature detection process.

20

CLAIMS

1. A method of detecting features-of-interest in an image, comprising the steps of:
5 making an intensity independent comparison of the image intensity profile with a shape model for predefined features of interest; and
storing for the profile a label comprising the level of similarity to the shape model.
- 10 2. A method according to claim 1 wherein the label further comprises a measure of the orientation of the feature.
3. A method according to claim 1 or 2 wherein the predefined features of interest comprise positive and negative step edges in the intensity profile.
- 15 4. A method according to claim 1, 2 or 3, further comprising performing the comparison on an intensity profile taken across successive frames of an image sequence to derive a measure of the velocity of image features moving through the sequence, and wherein said label further comprises said measure of the velocity.
- 20 5. A method according to claim 1, 2, 3, or 4, wherein the comparison comprises the steps of:
convolving the intensity profile with a quadrature filter to derive the phase and amplitude of a pair of odd and even components of the profile in the spatial or spatio
25 temporal frequency domains;
deriving from the difference between the components a measure of one of the feature symmetry or asymmetry along the profile;
deriving from the components a measure of the local amplitude along the profile;
30 defining areas of the image satisfying a predefined constraint on the value of said one of the feature symmetry or asymmetry as satisfying the shape model and thus

-18-

belonging to an image feature; and
wherein the label comprises said measures.

5 6. A method according to claim 5 wherein said quadrature filters are log-Gabor filters.

7. A method according to claim 5 or 6 wherein said predefined constraint on said one of the feature symmetry or asymmetry is that its magnitude is greater than or equal to a predefined threshold.

10 8. A method according to claim 5, 6, or 7, comprising the step of convolving the intensity profile of the image with quadrature filters oriented in different directions across the image to detect image features oriented differently in the image, and deriving from the relative responses of the filters a measure of the feature orientation.

15 9. A method according to any one of the preceding claims, further comprising tracking detected image features through successive frames of an image sequence by comparing the labels of areas in a succeeding frame with the labels of areas in the preceding frame over a predetermined search region to define as relating to the same feature those areas in each frame whose labels satisfy predefined search criteria.

20 10. A method according to claim 9, wherein the predetermined search region is defined as a region within a predetermined distance of the position of the detected image feature in an adjacent frame in the sequence.

25 11. A method according to claim 10, wherein the predetermined distance is set in dependence upon the velocity of the detected image feature.

30 12. A method according to claim 9, 10 or 11 wherein the predefined search criteria comprise a condition on the value of the feature symmetry or asymmetry of the areas in the search region.

-19-

13. A method according to claim 12 wherein the condition on the value of the feature symmetry or symmetry or asymmetry is that it is greater than a predetermined amount.

5 14. A method according to claim 12 or 13 wherein the predefined search criteria comprise the condition that the difference in the orientation of the detected features in the two frames is less than a predetermined amount.

10 15. A method according to any one of the preceding claims wherein the intensity independent comparison with the shape model is performed at a plurality of scales.

15 16. A method according to any one of the preceding claims wherein the comparison is performed on an intensity profile taken through a plurality of images of different depths of a volume being imaged.

17. A method according to any one of the preceding claims wherein the image is an ultrasound or magnetic resonance or x-ray or standard camera image.

20 18. A method according to any one of the preceding claims wherein the image is an ultrasound contrast image, an ultrasound Doppler image or an ultrasound harmonic image.

25 19. A method according to any one of the preceding claims wherein the image is a medical image.

20. A method according to any one of the preceding claims wherein the image is a remote sensing or non-invasive image.

30 21. A method according to any one of claims 1 to 16 wherein the image is an echocardiogram.

-20-

22. A method according to claims 20 wherein the feature is a ventricular or atrial wall.

23. A computer program comprising program code means for executing the
5 method of any one of the preceding claims.

24. An image processing system comprising image processing means adapted to perform the method of any one of the preceding claims.

Fig.1(A).




Fig.1(B).





Fig.2.









Fig.5(A).




Fig.5(B).



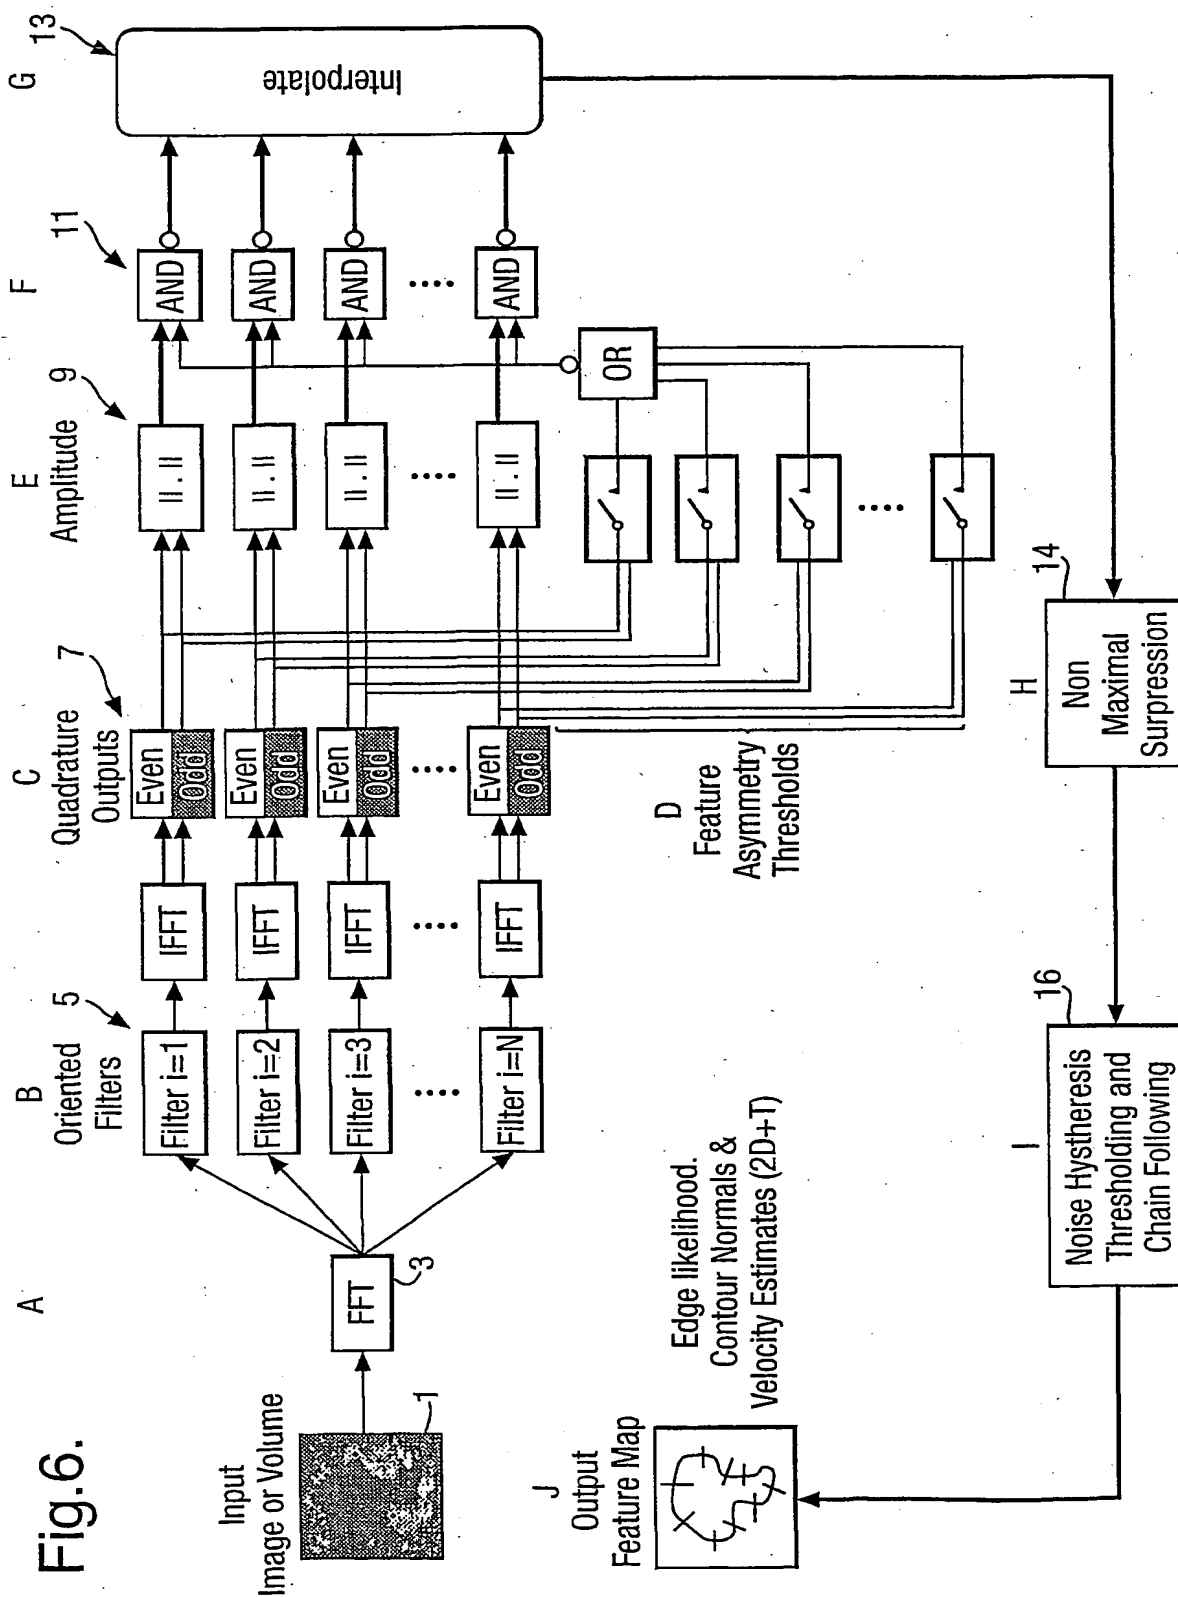


Fig.5(C).



4/7



5/7

Fig.7(A).




Fig.7(B).





Fig.8.

Scale=16



Scale=32



Scale=48



Scale=64



Scale=24



Scale=40




Scale=56



Scale=72



Fig. 9.



7/7

Fig.10.




Fig.11.

